

Data Analytics and Workflows Breakout

ASCR Workshop on Extreme Heterogeneity in HPC
Jan 23-25, 2018

Videoconference: <https://lbnl.zoom.us/j/121725281>

Slides: <https://goo.gl/2T6NPr>

Data Analytics and Workflows (DA&W) Contributors

Moderator(s):

Part 1: Tom Peterka, Shinjae Yoo

Part 2: Christine Sweeney, Wes Bethel

BOGists (46 on mailing list, ~20 on call):

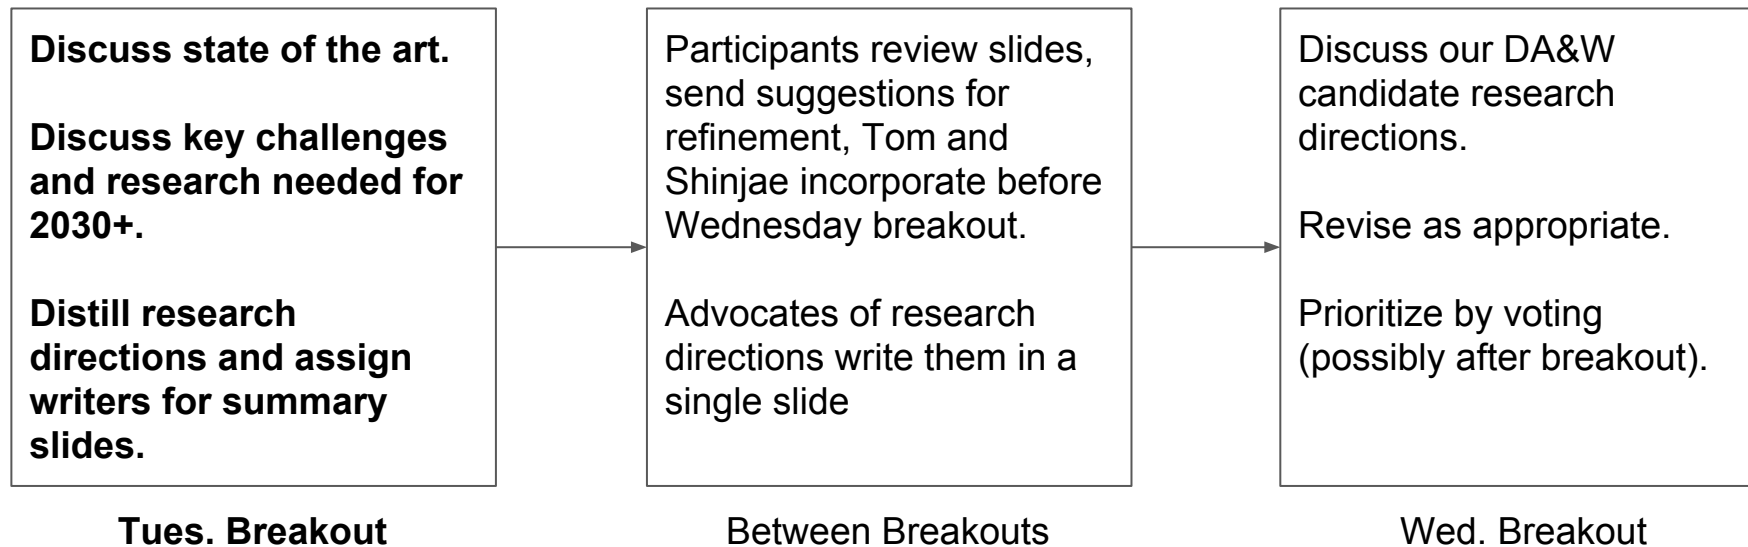
- Tom Peterka (tpeterka@mcs.anl.gov)
- Shinjae Yoo (sjyoo@bnl.gov)
- Christine Sweeney (cahrens@lanl.gov)
- Wes Bethel (ewbethel@lbl.gov)
- Kenneth Moreland (kmorel@sandia.gov)
- Jay Lofstead (nferrier@anl.gov)
- Andrew Younge (ajyoung@sandia.gov)
- Kevin Barker (kevin.barker@pnnl.gov)
- Ray Bair (rbair@anl.gov)
- Nicola Ferrier (nferrier@anl.gov)
- Adolffy Hoise (ahoise@bnl.gov)
- Andreas Gerstlauer (gerstl@ece.utexas.edu)
- Lavanya Ramakrishnan (lramakrishnan@lbl.gov)
- Stan Tomov (tomov@icl.utk.edu)
- Candace Culhane (culhane@lanl.gov)
- Line Pouchard (pouchard@bnl.gov)
- Brian Van Essen (vanessen1@llnl.gov)
- Ewa Deelman (deelman@isi.edu)
- Valerio Pascucci (pascucci@sci.utah.edu)
- Dave Pugmire (pugmire@ornl.gov)
- Peter Kogge (kogge@nd.edu)
- Larry Kaplan (lkaplan@cray.com) (day 2 only)

Reminder of Our Charge

The purpose of this workshop is to identify the priority research directions for ASCR in providing a smart software stack that includes techniques, such as deep learning to make future computers composed of a variety of complex processors, new interconnects and deep memory hierarchies easily used by a broad community of computational scientists.

DA&W Breakout: The Plan

Goal for this breakout: Identify ~5 possible/promising research directions that address key challenges for DOE mission in the 2030+ timeframe. Focus on aspects related to heterogeneity.



Goal: DA&W Capability Targets and Research Directions

- **What DA&W capabilities will be needed in the 2030+ timeframe to make productive use of increasingly heterogeneous systems for DOE mission?**
- **What research is required to get from where the capabilities are now to where to where they need to be by 2030?**

Try to keep these at a high level (think topics in a call).

Process

- Identify current capability
- Identify target capability
- Identify challenges (i.e., $\text{gaps} = \text{target} - \text{current capabilities}$)
- List research directions needed to solve challenges (possibly part 2)
- Prioritize research directions (part 2)

Current and Target Capabilities

Current Capabilities 1/2

- Hetero workloads (a.k.a. use cases)
 - Big compute (HPC)
 - In situ HPC visualization and analysis workflows
 - Big data
 - Distributed, cloud, web workflows
 - High Throughput Computing (HTC)
 - Post hoc visualization and analysis tools
 - Experiment, observation
 - Integration
 - Starting to combine
 - Time constraints, reproducibility unknown or unenforced
 - Usability and productivity
 - Increasingly complex interfaces / tools
 - Performance & Portability
 - Node / system wide
 - VTKM, TensorFlow, Kokkos, others support CPU/GPU
 - MPI+X paradigm
 - Note for target -> node / system wide (CPU/GPU/FPGA/NC/QC...)
 - Related to reproducibility as nodes increase, it will impact reproducibility

Current Capabilities 2/2

- Heterogeneous hardware
 - HPC
 - Even not all HPC systems are the same
 - Clouds
 - Integration
 - Currently some portability, but not performance portability
- Hetero software stacks
 - HPC
 - Apache Hadoop-like big data
 - Integration of HPC + Big Data
 - Poor performance now
 - Limited examples of Spark + MPI
 - Containerization in HPC (i.e., TensorFlow in one node)
 - Interfaces are heterogeneous too
- Algorithms (analysis)
 - Mostly homogeneous, deterministic data reduction
 - VTK-M tries to address some of these issues
 - Not much approximate, stochastic algorithms targeting new architectures
- Data models
 - Representations derived from multidim arrays
 - Generalized programming models and workflow composition languages

Target Capabilities 1/4

- Hetero workloads
 - Fully integrated HPC, big data, experimental, observational workflows
 - E.g., Tight integration of experiments and simulations
 - Full support of streaming/real-time processing
 - Data analytics of multiple modality / data assimilation (data fusion)
 - Future data infrastructure support for hetero h/w (need more brainstorming)
 - Getting data to/from quantum and neuro computers
 - Unstable to read/write
 - Memory part of the compute (not Von Neumann)
 - Efficient data movement over hetero memory architecture (i.e. architecture aware)
 - Workflow exposure to or abstraction from architecture details

Target Capabilities 2/4

- Hetero hardware
 - Algorithms for existing hetero h/w
 - Performance-portable
 - Algorithm developments for future hardware
 - Stochastic and unstable
 - neuro, quantum
- Hetero software stacks
 - Interoperable and coexistent software stacks
 - E.g., big data and HPC

Target Capabilities 3/4

- Algorithms
 - Data fidelity -- reproducibility
 - Approximation uncertainty varies by platform
 - Scientific validity
 - Data exploration and search
 - Steering an objective function
 - Reproducibility, accuracy, optimization
 - Spectrum of domain expert knowledge <-> hybrid <-> ML data driven
 - Multimodality (multiscale, multisource)
 - Usability and programmability
 - Hiding h/w and s/w complexity (i.e. VTK-M)
 - Enabling interoperability and coexistence
 - Using future programming models to improve usability
 - Driven by users (not developers)
 - Performance
 - Measuring, predicting, enforcing, performance of complex workflows
 - Autonomic control (e.g, auto-tuning using AI)

Target Capabilities 4/4

- Data models
 - Unstructured objects
 - Approximate data model
 - Opportunity for data cleaning, smoothing
 - Variable precision (i.e. (sub)-single precision)
 - Map to computing environment
- Workload Management Tools
 - Abstraction
 - Interaction between the layers
 - Software infrastructure
 - Resource Allocation
 - Managing workflows when resources are not allocated monolithically or statically
- Interpretability of deep learning
 - Understand / explain the reason AI model works

Research Challenges and Opportunities

Research Challenges and Opportunities 1/2

- Complex/hierarchical workflow profiling and provenance (Line Pouchard)
- Dynamic workflows (Christine Sweeney & Tom Peterka)
 - Migrate around system depending on resource availability
- Data algorithms (Tom Peterka & Shinjae Yoo)
- Data Models (Dave Pugmire)

Research Challenges and Opportunities 2/2

- Rapid development of analytics for multiple modality (Brian Van Essen)
- Resource isolation and sharing (Jay Lofstead & Andrew Younge)
- Usability of resource allocation (Ken Moreland & Lavanya Ramakrishnan)
 - Some form of predicted model is required
- Software infrastructure interoperability and coexistence (Ewa Deelman & Tom Peterka)
- Performance Portability (Valerio Pascucci)

Overview of Data Analysis and In Situ Workflow (BOG 3) Priority Research Directions

Data Analytics and Workflows (DA&W) Contributors

Moderator(s):

Part 1: Tom Peterka, Shinjae Yoo

Part 2: Christine Sweeney, Wes Bethel

BOGists (46 on mailing list, ~20 on call):

- Tom Peterka (tpeterka@mcs.anl.gov)
- Shinjae Yoo (sjyoo@bnl.gov)
- Christine Sweeney (cahrens@lanl.gov)
- Wes Bethel (ewbethel@lbl.gov)
- Kenneth Moreland (kmorel@sandia.gov)
- Jay Lofstead (nferrier@anl.gov)
- Andrew Younge (ajyoung@sandia.gov)
- Kevin Barker (kevin.barker@pnnl.gov)
- Ray Bair (rbair@anl.gov)
- Nicola Ferrier (nferrier@anl.gov)
- Adolfy Hoise (ahoise@bnl.gov)
- Andreas Gerstlauer (gerstl@ece.utexas.edu)
- Lavanya Ramakrishnan (lramakrishnan@lbl.gov)
- Stan Tomov (tomov@icl.utk.edu)
- Candace Culhane (culhane@lanl.gov)
- Line Pouchard (pouchard@bnl.gov)
- Brian Van Essen (vanessen1@llnl.gov)
- Ewa Deelman (deelman@isi.edu)
- Valerio Pascucci (pascucci@sci.utah.edu)
- Dave Pugmire (pugmire@ornl.gov)
- Peter Kogge (kogge@nd.edu)
- Larry Kaplan (lkaplan@cray.com) (day 2 only)

View from 15 years out:

- Reduce complexity, increase speed of formulating complex scientific analytics processes suitable for use on EH platforms, use EH software
- Draw from heterogeneous sources of analytics tools and systems, deploy on heterogeneous platforms, adapt to changing conditions, be reproducible.
- User view: specify compound analytics processing with high-level descriptions of actions.
- These are reduced to use of specific tools, with specific parameters, that best match common motifs drawn from ethnographic studies/best practices, as well as previous “recipes” built from other users, with recommendations from AI systems that suggest good combinations of tool/(hyper-)parameter/platform to meet some specific performance objectives.
- The tools and components should themselves be capable of running effectively on EH platforms, which include increasingly complex memory/storage hierarchies/services for resource management.
- These tools may come from completely disparate sources, have disparate execution models.

Challenges

Challenges:

Data heterogeneity: data in many formats/forms, coming from many sources, sinks, shared by different workflow infrastructures and programming models, partial or approximate.

Heterogeneous dynamic software environments: workflows that can interface and integrate with each other, workflows that are complex and dynamic, need for understanding performance and for reproducibility despite complexity and dynamicism

Algorithms: algorithms that can work with uncertainty and approximation, utilize data encountered (multiscale, multimodal) , be ready with the right algorithms at the right time (data-driven, domain-specific), agile creation of algorithms.

Usability / Productivity: cross-cutting challenge that hits at all levels

An Integrated View of the PRDs

Stacked	Cross-cutting
Automated Data Analytics Algorithms PRD 3.1	Portable productivity and Performance PRD 3.5
Integrated Dynamic Software Environments PRD 3.2, 3.3	
Data Model Abstractions PRD 3.4	

PRD 3.1

Rapid, Semi-autonomous Data Analytics Algorithm Development

B. Van Essen, T. Peterka, S. Yoo

Data analytics workflows on heterogeneous systems will require specialized algorithms for in situ analysis of multiple data modalities. The development of specialized machine learning / data analysis algorithms will need to be automated for scientific application by domain scientists, not just data scientists. New algorithms will need to exploit heterogeneous hardware such as new storage and computing models. In the face of such heterogeneity, results will vary in fidelity and reproducibility, and algorithms will need to incorporate that uncertainty in their methods and results.

PRD 3.2

Analytics Ecosystems Programming and Usage

C. Sweeney, K. Moreland, L. Ramakrishnan, E. Deelman, T. Peterka

Future heterogeneous workloads will feature integrated computational, big data, experimental, observational, streaming, real-time processing, and multimodal data assimilation. No single data ecosystem will support all these uses while fully exploiting heterogeneous architecture. EH systems will need to accommodate and integrate complex, dynamic, and disparate workflows and software stacks. Workflow execution needs to be agile, not monolithic and predetermined; and workflows need to dynamically manage conflicting resource demands and balance work on heterogeneous resources.

PRD 3.3

Analytics Ecosystems Validation

L. Pouchard, S. Yoo

Workflows are used for in situ streaming and analysis workflows, but without profiling and provenance of workflows, scientific code optimization for performance and reproducibility will be left to chance. Detailed workflow profiling and the provenance of both scientific results and performance is essential to extract metrics enabling code comparison, optimization, and reproducibility. Scientists will be able to determine the impact of various code parameters on reproducibility, performance, resource usage, and understand the interplay of different architecture designs and their workloads.

PRD 3.4

Data Model Abstractions

D. Pugmire, W. Bethel

Data Models act as abstractions of information that overlay the underlying data. These abstractions need to support a wide range of areas related to data, operations on the data, and how operations are scheduled and executed in a heterogeneous system. Efficient use of the EH system demands good abstractions to deal with the complexity. More flexible and adaptable data models are needed to support an increasingly heterogeneous analytics software landscape that integrates software methods and implementations from many different sources.

PRD 3.5

Data Analytics Portability, Productivity, and Usability

V. Pascucci, L. Ramakrishnan

With heterogeneous architectures, the challenge of porting and scaling performance of data analytics algorithms, workflows, and data models across target resources will increase exponentially. Performance portability, therefore, is key to data science. Increased software and hardware complexity of future systems will also impede usability and productivity of computing resources. If we do not address usability and productivity in the data ecosystem, it will be difficult if not impossible to achieve the needed efficiency gains on future heterogeneous systems.

PRD Details

Potential Research Directions

- PRD 3.1.1 Rapid, Semi-autonomous Development of Multimodal Data Analytics
- PRD 3.1.2 Evolving Data Algorithms for the EH Era
- PRD 3.2.1 Resource Allocation Usability
- PRD 3.2.2 Software Infrastructure Interoperability and Coexistence
- PRD 3.2.3 Dynamic Workflows
- PRD 3.3 Complex Workflow Profiling and Provenance
- PRD 3.4 Data Models
- PRD 3.5.1 Performance Portability
- PRD 3.5.2 Address usability and productivity in the software stack for next-generation heterogeneous infrastructure
- (move to OS/R) Resource Isolation and Sharing

PRD 3.1.1 Rapid, Semi-autonomous Development of Multimodal Data Analytics (B. Van Essen)

Data analytics and adaptive workflows on heterogeneous systems will require specialized machine learning algorithms to provide in-situ fusion of multiple data modalities. These specialized machine learning / deep learning (ML/DL) algorithms will need to be rapidly developed for each scientific application and ideally can be developed by domain scientists, not just ML experts. The goal should be that a domain scientist can specify the problem of interest, identify the key data types to be evaluated, and describe the metric used to evaluate the goodness of a solution (or distance between two points in the problem space). Using these inputs an advanced ML/DL system would be able to create an optimized data analytic for this heterogeneous scientific application.

- Research challenges/opportunities: automatic exploration and tuning of ML models (network topologies and hyper-parameters), abstract ML / DL specification interface, portable ML / DL model descriptions
 - Metrics for progress: standardized tools for exploring / tuning ML / DL model, guided ML/DL model search, automated ML / DL development
 - How to specify the problem statement at a high level, and how to guarantee it is being translated correctly into combinations of tools and their appropriate parameter settings.
- Potential research approaches and research directions: ML / DL tool for multi-model datasets, tools for automatic ML / DL model exploration, tools for tuning ML / DL models, creating high level interface for specifying ML / DL goals (e.g. input data, objective functions / metrics)
- How and when will success impact technology? Each stage of progress on this topic will incrementally improve the process for creating new ML / DL models on new problems.
- Why make this a priority research direction? Advanced machine learning and deep learning are one of the most promising technologies for fully exploiting the capabilities of an extremely heterogeneous system, but they are extraordinarily tricky to develop and must be optimized for each scientific application. Automating the development of these data analytics is critical to empowering enable domain scientists to tackle new problems.

PRD 3.1.2 Evolving Data Algorithms for the EH Era (T. Peterka & S. Yoo)

Heterogeneous hardware and workloads present opportunities and challenges for data analytics algorithm development. New algorithms will be needed to exploit heterogeneous hardware such as deep memory/storage locality and new compute devices. In the face of such heterogeneity, performance and scientific results will vary in fidelity and reproducibility, and algorithms will need to incorporate that uncertainty in their computations.

- Research challenges: Data fidelity and reproducibility where approximation uncertainty varies by platform; multimodal (multiscale, multisource) data; usability and programmability with future programming models; hybrid data-driven and domain-specific algorithms for science
- Metrics for progress: new metrics needed for performance, usability, validity
- Potential research approaches and research directions: approximate algorithms, scalable parallel deep learning algorithms, online hyper- and parameter search, uncertainty quantification, soft error bounds
- How and when will success impact technology: New algorithms can capitalize on new h/w capabilities, just as algorithms have been rewritten for GPUs
- Why make this a priority research direction: Old algorithms on new architectures won't run or at best won't scale, and the results will not be meaningful.

PRD 3.2.1 Resource Allocation Usability (K. Moreland & L. Ramakrishnan)

- Description
 - Current systems make no effort to manage heterogeneous resources
 - EH systems require apps to query/request own resources from the workflow.
 - Workflow needs to dynamically manage these heterogeneous resources.
- Research challenges
 - Workflow must manage many different, conflicting, resources (e.g. MPI vs PGAS vs AMT, OpenMP vs pthreads vs C++11 concurrency vs TBB vs HPX)
 - Workflow must dynamically balance work on heterogeneous resources
 - Workflow must dynamically reconfigure resources
- How and when will success impact technology?
 - Technology could help today. Now struggle with configuring resources
- Why make this a priority research direction?
 - As heterogeneity increases, managing resources through a job scheduler becomes unmanageable
 - Required for workflow managing components on heterogeneous resources

PRD 3.2.2 Software Infrastructure Interoperability and Coexistence (E. Deelman & T. Peterka)

Heterogeneous workloads will feature fully integrated HPC, big data, experimental, observational workflows, support streaming and real-time processing, and multimodal data analytics and assimilation. No single data infrastructure will be able to support all these and future unforeseen use cases while simultaneously exploiting the opportunities afforded by heterogeneous hardware platforms. Hence interoperability between coexisting software infrastructures is needed.

- Research challenges: disparate programming models, incompatible execution models, hetero OS/R system services (scheduling, allocation, security, storage)
 - Metrics for progress: basic functionality -> performance -> usability
- Potential research approaches and research directions: programming model interface adapters, execution launchers, system service abstractions
- How and when will success impact technology: Software interoperability is needed now and will continue for the next 15 years. Each software stack has its own strengths, and research in integration can capitalize on those strengths without reinventing technology.
- Why make this a priority research direction: Economy. Software interoperability is the most economical way to reuse software from industry and commerce. Just as we rely on vendors to build hardware, we are using heterogeneous software stacks from multiple domains now and will continue to do so.

PRD 3.2.3 Dynamic Workflows (C. Sweeney & T. Peterka)

EH systems with heterogeneous elements will need to accommodate complex and dynamic workflows including data analytics workflows. Data analytics, largely dependent on underlying software stack, interfaces to parallelism, I/O and OS, will need to become more agile and less monolithic or pre-determined. Data coming from various sources at varying rates will contribute to the dynamic workflow.

- **Research challenges**
 - Effective and often efficient data analytics workflows despite uncertain data location, timing, processing resources, and partial data.
 - Take advantage of where data is and what processing is available to fit in useful analytics where possible.
 - Metrics for success: high utilization of EH systems for data analytics where workflow progresses and is effective
- **Potential research approaches and research directions**
 - “Assessment” capabilities to identify analysis needs/opportunities dynamically
 - Workflow infrastructure (building blocks of a workflow system) and service-oriented workflow paradigms.
 - Probabilistic, approximate, incremental, real-time, streaming, multi-streaming data analytics algorithms, including ML.
 - Visualization and analytics of status and performance of complex and dynamic workflows within the EH system.
 - Provenance that allows dynamic workflows to be reproduced or understand how results were obtained.
 - Data analytics workflows that span both EH systems and across wide area network
- **How and when will success impact technology?**
 - For new dynamic data analysis algorithms, some benefit could be seen immediately.
 - Viz and analysis of the EH system could begin and benefit as soon as these new components become available.
 - For infrastructure changes or new software paradigms, speed of change will depend on productivity of interfaces to EH
- **Why make this a priority research direction?**
 - Dynamic workflows will be present in both data analytics and applications
 - It advances knowledge in data analytics algorithms and workflows
 - Makes better use of EH systems, provides understanding of the systems and security.

PRD 3.3 Complex Workflow Profiling and Provenance (L. Pouchard)

- Workflows can be used for in-situ streaming and analysis applications coupled to numerical simulation codes
 - they enable scientists to reduce, analyze and prioritize their scientific results as they are created
 - they can be a way to orchestrate the interaction between e.g. multi-level physics codes.
 - detailed workflow profiling and the provenance of both scientific results and performance extract metrics enabling code comparison and optimization
- Research challenges:
 - which metrics and at what granularity need extracting to enable code comparison and optimization for complex workflows on EH architectures?
 - determining tradeoff levels and their impact on reproducibility vs. code performance vs. resource usage for complex workflows on EH?
 - Progress: understanding which architecture design emphasizes which of the aspects above
- Research approaches and directions:
 - tools to extract required metrics for workflows from EH components and devices
 - quantification of acceptability thresholds given emphasis on either aspect per application community
 - design profiling workflow campaigns and build training sets to create ML profiling models
- Without profiling and provenance of workflows scientific code optimization for performance and reproducibility will be left to chance and not progress in a systematic manner

PRD 3.4 Data Models (D. Pugmire & W. Bethel)

Data Models act as abstractions of information that overlay the underlying data. These abstractions need to support a wide range of areas related to 1) data, 2) operations on the data, and 3) how operations are scheduled and executed in a heterogeneous system.

Research challenges:

- Abstractions of data that include: location in storage hierarchy, provenance, reductions, quality, precision, accuracy
- Abstractions that allow mapping operations on data onto particular hardware, decomposition of analysis into sub-tasks (and recombination), respecting resource limits (space and time) for tasks
- Understanding tradeoffs between underlying representations of data and the wide variety of available hardware
- Finding paths forward that allow for using and combining heterogeneous methods, from diverse sources, to meet DOE science mission needs

Potential approaches: Identification of classes/motifs for data and operations on data. Prediction models for performance as a function of the underlying data. Models for error bounds and error propagation as a function of the underlying data.

How and when will success impact technology: Abstractions will allow for efficient execution (either manually or via a workflow system) of operations on data throughout the system.

Why make this a priority research direction: Efficient use of EH system not likely without good abstractions to deal with the complexity. More flexible and adaptable data models are needed to support an increasingly heterogeneous (analytics) software landscape, where we want to leverage and combine s/w methods/implementations from many different sources (industry, research, etc.)

PRD 3.5.1 Performance Portability (V. Pascucci)

With heterogeneous architectures the challenge of porting and scaling performance of data analytics and workflows across target resources will be exponentially harder.

- Research challenges:
 - Provide high level descriptions of analytics and workflows that are unbiased with respect to particular platforms (e.g. data movements vs compute)
 - Develop runtimes that implements analytics pipelines on diverse hardware resources
 - Interoperate with diverse data sources (simulations or experiments)
 - Maintain possible needs for user interaction, streaming, ...
- Success metrics: Efficient execution on diverse platforms, scalability of high level descriptions of analytics workflows ~~pipelines~~, user productivity, percentage of peak performance achieved.
- Possible approaches: Programming models and runtime systems with unique needs for analytics and workflows such as direct user interaction
- Success will have immediate impact since this is already a problem for platforms and software stacks.
- This is a high priority direction because it is key to executing science workflows and understanding their results

PRD 3.5.2 Address usability and productivity in the software stack for next-generation heterogeneous infrastructure (K. Moreland and L. Ramakrishnan)

The complexity of future heterogeneous infrastructure will further increase the current challenges with usability and productivity of using software tools and hardware infrastructure.

Research challenges:

- Need a understanding of user workflows and their interaction of heterogeneous infrastructure to develop a productivity metric
- Ability to seamlessly integrate and dynamically manage heterogeneous resources in user workflows
- Balancing usability and productivity with performance and portability trade-offs.

Potential research approaches and research directions

- Use of ethnography and user research to understand the productivity metrics
- Development of new methods and interfaces to manage heterogeneous resources in user workflows
- Balancing usability and productivity with performance and portability trade-offs across the software stack
- Consider performance models for better prediction.

How and when will success impact technology? While short-term impact will improve usability of resources, key research will need deeper understanding and interaction across the stack and development of appropriate metrics which will take 8-10 years.

Why make this a priority research direction? If we do not address usability and productivity in the software stack, it will be difficult if not impossible to achieve the efficiency gains that will need to be achieved on next-generation systems.

Backup Slides

Additional (Orphan) Research Challenges / Opportunities

- Scalability of workflow tools and integrability with other tools
- Secure computation
- Integrated / encompassing / interoperable workflows
- Interplay of system component
 - OS Runtime
 - Interface of components
 - S/W components: see D Pugmires Data Model PRD slide
 - H/W or other system components: may provide services needed by analytics workflows, there may be heterogeneity in how these services are accessed, requested, etc.

Possible Priorities: Choosing the best PRDs

- “Usable, useful, secure”
- advances knowledge in that area
- is competitive with other research
- balances innovation with risk of failure
- optimizes scientific return for resources
- maximizes previous investments
- maximizes currently available technology or technology TBA
- is in line with scientific goals in 2030 (where we think applications will be)
 - i.e., looks far enough forward (10-15 years)
- is unique to data analytics and workflows
- solves an *EH* challenge or capitalizes on an EH opportunity (is specific to EH)
- extends something we are already good at
- does not duplicate research directions we see happening elsewhere

PRD 3.2 Resource Isolation and Sharing (Jay Lofstead and Andrew Younge)

- Research challenges (EH?)
 - Exascale compute will feature both large scale, single application runs as well as composite applications and workflows. The latter case is what we are focused on, as it will be increasingly important for heterogeneous architectures. In these cases, we want to efficiently share data between components while maintaining isolation and separate resilience domains. We need OS support for on node isolation and data sharing as well as techniques that require little to no different programming models to compose workflows off node. Currently, we have demonstrated in the Hobbes project the ability to share data between isolated partitions on a single node, however there is a need to address this for programming and general use cases.
 - Techniques and technology for co-hosting isolated components on a single node need to be developed. Currently we have rudimentary technology for partial solutions, such as containers or virtual machines, but they do not have good technology or techniques to offer explicit, controlled, and deliberate sharing within a single node. No specific mechanism exists for inter-node use.
- Potential research approaches and research directions
 - Efficient on node isolation and co-hosting
 - Techniques for sharing between components on a single node securely and without coupling resilience domains
 - Techniques for extending the above to inter-node models
 - Communication and consistency protocols and approaches to manage the parallel N-M connections to ensure that processing should proceed with a validated (and complete) data set.
- How and when will success impact technology?
 - This technology can be rolled out with operating system and programming library deployment. This implies that it can be rolled out shortly after it is hardened for production use.
- Why make this a priority research direction?
 - Extreme-scale machines and workloads will require simultaneously deploying multiple different software stacks to most efficiently handle their own processing. This will become necessary as different workloads must leverage desperate heterogeneous capabilities at different rates.
 - This is both for separate jobs as well as for a single composed application set or workflow. Only through techniques and technologies like these will we be able to efficiently deploy large scale, diverse workloads on our extreme scale machines.